

Description

Totem-Pole Power Control for Processors

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation in part application of a provisional patent application of the same name, serial number 60/479,707, filed 19 June, 2003.

BACKGROUND OF INVENTION

[0002] Prior art power controllers for processors are usually buck converters, as shown in figure 1. They suffer from poor response. Multi-phase buck converters, as shown in figure 2, are a significant improvement, but they are still slow.

[0003] A switched-current power converter, as taught by US Patent No. 6,121,761, "Fast Transition Power Supply", issued 19 September, 2000, and shown (simplified) in figure 3 is much faster. A variant (US Patent Application serial number. 10/709,484, "Switched-current Power Converter" filed 8 May, 2004) having a plurality of switched-current sources is shown in figure 4. The plurality of cur-

rent sources can be generated using one current source and a matrix transformer. The current can transition from zero to full load and back to zero very quickly, essentially the time that it takes to close, then open, solid-state switches. The above patent and patent application are owned by the same entity as the present invention, and they are incorporated herein by reference.

[0004] In the power converter of figure 4, the switches are located outside of the processor, though preferably they are very close to it.

SUMMARY OF INVENTION

[0005] Figure 5 shows that for optimum timing and minimum loop inductance, the switches should be located within the processor, either immediately proximate to the processor die or embedded within it. If the switches are shown as MOSFETs, the circuit of figure 6 is the result. Note, the voltage on the input will be low or high, respectively, depending upon whether the bottom or the top MOSFET switch is turned on.

[0006] Figure 7 shows a totem pole output driver. Note the similarity to the switched-current converter cell of figure 6. When the bottom MOSFET is on, the output is low, when the top MOSFET is on, the output is high, with a voltage

somewhat lower than the V_{cc} bus. The waveform is shown in figure 8.

[0007] Figure 9 shows that a constant current can be injected into a totem pole cell, and the voltage will be low if the bottom MOSFET is on, and somewhat over V_{cc} if the top MOSFET is on. The waveform is shown in figure 10. Assuming a relatively high impedance elsewhere on the data output line, the current will either flow through the top MOSFET to V_{cc} or it will flow through the bottom MOSFET to the return. If a plurality of constant current sources are connected to a plurality of totem pole cells, the current flowing into V_{cc} can be any amount from zero to the sum of the current sources, and it can transition as fast as the totem pole cells can change state.

[0008] A totem pole cell may be optimized for power control or as a data driver, depending upon the amount of current it controls. A larger, higher current totem pole cell may be too slow for high speed data flow, but none the less it may be able to handle some data. A smaller, faster totem pole cell may handle too little power to be characterized as a power controller, but non-the-less, its contribution to power control may be beneficial and collectively with many other similar totem pole cells, may be significant.

BRIEF DESCRIPTION OF DRAWINGS

- [0009] Figure 1 shows a prior art buck converter.
- [0010] Figure 2 shows a prior art multi-phase buck converter.
- [0011] Figure 3 shows a single section of a prior art switched current power converter.
- [0012] Figure 4 shows a prior art switched current power converter.
- [0013] Figure 5 shows a single section switched current power converter in which the current switches are inside a processor package.
- [0014] Figure 6 shows that a practical switch arrangement for a section of a switched current power converter within a processor package bears a strong resemblance to a totem pole bus driver circuit.
- [0015] Figure 7 shows a prior art totem pole data driver. Figure 8 shows its voltage waveform.
- [0016] Figure 9 shows a single section of a switched current power converter that can also function as a data driver. Figure 10 shows its voltage waveform.
- [0017] Figure 11 shows a processor having a plurality of combined switched current power converter inputs and data outputs driving a data bus to a data receiver. A ballast

block divides a current from a source of current into a plurality of generally equal smaller currents.

[0018] Figures 12, 12a and 12b show a modified "NRZ" data waveform.

[0019] Figures 13a, 13b, 14a and 14b show a modified "NRZ" data waveform further having a variable duty cycle.

[0020] Figure 15 shows an example of a processor with a switched power converter with a variety of current and voltage control means, to show the flexibility of this invention when combined with a switched current power converter.

DETAILED DESCRIPTION

[0021] Prior art power controllers for processors frequently are buck converters, for example, the buck converter 1 shown in figure 1. They suffer from poor dynamic response. An inductor 2 and a capacitor 5 average the voltage output V_o to equal the input voltage V_i times the duty cycle of a switching means 3. A catch diode means 4 conducts current into the inductor 2 when the switching means 3 is open.

[0022] Multi-phase buck converters, for example, the multi-phase buck converter 10 shown in figure 2, are a significant improvement, but they are still slow. A plurality of

inductors 11 through 14 and a capacitor 23 average the output voltage V_o to equal the input voltage V_i times the duty cycle of a plurality of switching means 15 through 18. A plurality of catch diodes 19 through 22 conduct current into the respective inductors 11 through 14 when the respective switching means 15 through 18 are open.

[0023] A switched-current power converter 30, as taught by US Patent No. 6,121,761, Fast Transition Power Supply, a single section of which is shown in figure 3, is much faster. A current I from a source of constant current 31 is directed by a switching means 32 either to return or to a capacitor 33 to supply a current I_o to a voltage source input V_{cc} of a processor 34. A controlled current I_c into the capacitor 33 can be zero or equal to the current I , depending upon the state of the switching means 32, and can transition between values as fast as the switching means 32 can switch. This can be very fast indeed. The single section is shown for illustration and to explain the principle of operation. A practical switched current power converter has a plurality of similar sections, as explained below.

[0024] Figure 4 shows a prior art switched current power converter 40 (US Patent Application No. 10/709,484, "Switched-current Power Converter") having a plurality of

–current sources 41 through 45 each producing an essentially equal constant current I . The plurality of current sources 41 through 45 can, as an example, not a limitation, be generated using one current source and a matrix transformer. The current can transition from zero to full load and back to zero very quickly, essentially the time that it takes to close, then open, a plurality of solid-state switching means 46 through 50.

[0025] In the switch current power converter 40 of figure 4, the plurality of switching means 41 through 50 are located outside of a processor 52, though preferably they are very close to it. In operation, the plurality of switching means 46 through 50 may switch in various combinations to switch currents from the respective current sources 41 through 45 individually either to ground or to a capacitor 51. A control current I_c equals the number of the plurality of switches 46 through 50 that are switched to the capacitor 51 times the current I . The value of the control current I_c can vary as rapidly as the plurality of switching means 46 through 50 can change state. With solid state switches, like MOSFETs as an example, not a limitation, that can be very fast indeed, nearly instantaneous.

[0026] Figure 5 shows a single stage switched current power

converter 50 arranged for optimum timing and minimum loop inductance by having the switching means 63 located within the processor package 62, either immediately proximate to the processor die or embedded within it. A capacitor 64 may be outside the processor package 64, as shown, or part or all of the capacitance may be inside the processor package 64. Most modern processors have some internal capacitance, but the bulk of the capacitance is usually external as shown. The capacitor 64 stores charge, and to the extent that the charge on the capacitor 64 can be held constant, the voltage V_{cc} will be constant. Holding the charge constant, and thus the voltage V_{cc} , is the objective of controlling the current I_c very rapidly to conform, on average, to the output current I_o .

[0027] If the switching means 63 is shown as MOSFETs, the switched current power converter 70 of figure 6 is the result. Note, the voltage on the input will be low or high, respectively, depending upon whether the bottom MOSFET switch 74 or the top MOSFET switch 73 is turned on. An external capacitor 64 is shown, to store charge.

[0028] In figures 5 and 6, single sections are shown to explain the function of the switched power converter, but in a practical system a plurality of similar circuits would be

used, as shown in figure 4.

[0029] Figure 7 shows a prior art totem pole output driver circuit 80. In a processor 81, if a first MOSFET 82 is turned on, a data output signal 84 is "HIGH", with a voltage somewhat lower than the V_{cc} bus. If a second MOSFET 83 is turned on, the data output signal 84 is "LOW". Note the similarity to the switched-current converter cell of figure 6. The waveform is shown in figure 8.

[0030] Figure 9 shows that in a switched current converter 90 a constant current I from a source of constant current 94 can be injected into a totem pole cell of a processor 91, and the voltage will be low if the bottom a MOSFET 93 is on, and somewhat over V_{cc} if a top MOSFET 92 is on. The waveform is shown in figure 10. Assuming a relatively high impedance elsewhere on the data output line, the current will either flow through the top MOSFET 92 to V_{cc} or it will flow through the bottom MOSFET 93 to the return. If a plurality of constant current sources are connected to a plurality of totem pole cells, the current flowing into V_{cc} can be any amount from zero to the sum of the current sources, and it can transition as fast as the totem pole cells can change state. Thus a plurality of totem pole cell like circuits can be used in reverse to con-

trol the voltage on V_{cc} just as in the switched power converter 40 of figure 4.

[0031] A totem pole cell with current injection from an external constant current source may be optimized for power control or for data, depending upon the amount of current it controls. A larger, higher current totem pole cell may be too slow for high speed data flow, but none the less it may be able to handle some data. A smaller, faster totem pole cell may handle too little power to be characterized as a power controller, but non-the-less, its contribution to power control may be beneficial and, collectively with many other similar totem pole cells, may be significant.

[0032] In addition, the totem pole cell of figure 9 which sinks the current from a constant current source in either state may be a superior data cell, having a larger voltage margin and thus a greater noise immunity. Whereas a top MOSFET with higher resistance degrades performance when it causes a voltage drop from V_{cc} , it may benefit a current sinking cell, as it would increase the high state voltage. Assuming a wide data bus with a large number of data cells which sink current in the high state, the contribution to V_{cc} is twofold. First, when high, it is not a load on V_{cc} , as in a usual totem pole driver. Second, when high, it is a

source of current into Vcc, therefore reducing the current required from other sources, such as a VRM Pod, for processor functions..

[0033] A totem pole cell that sinks a constant current source can be a tri-state driver only if there is some provision for the current to continue to flow during the off state. If it is a bi-directional data bus, one end of the bus or the other must always be on, to provide a current path, and both ends can be on, overlapping, during transition as long as both are in the same state (HIGH or LOW). Another possibility would be to remove the drive currents during the off state, particularly for long quiescent times.

[0034] It is practical to generate a plurality of generally equal constant current sources with a matrix transformer as taught in 10/709,484. If a matrix transformer is operated at 100 percent duty cycle, so that one side of the push pull winding or the other is always on, or if a full bridge or half bridge circuit has current flow in one direction or the other without an off time, and if the transformer is powered by a constant current source, then the rectified output of every stage is a constant current source as well. With a one turn primary and a one turn secondary in each stage, the rectified outputs will comprise a plurality of

constant current sources with each one equal to the primary input current.

[0035] On the other hand, generating a large number of small, generally equal constant current sources is more of a problem. Figure 11 shows a processor 101 having a plurality of totem pole data cells 102-102 that send data to a data receiver 103 having a plurality high impedance data sense inputs 104-104. A current I from a source of constant current 105 can be divided into a plurality of generally equal constant currents using a ballast block 106. If currents are directed into a number of parallel paths, the currents will divide equally if the average voltage and the impedance of each paths are equal. Therefore, with some care, the current I from the source of constant current 105 will divide fairly well if the resistances of the several data bus lines are equal and the impedances of the MOSFETs comprising the plurality of totem pole data cells 102-102 are equal, provided that the totem pole data cells 102-102 are all in the same state. (Intentionally unequal currents can be generated as well, by intentionally making the impedances unequal).

[0036] The ballast block 106 may be a small block of magnetic material, perhaps ferrite or powdered metal, as illustra-

tions, not limitations, with through holes having wires in them comprising a plurality of very small parallel inductors. Small, so that they are inexpensive and take up very little board real estate, but still large enough so that the currents through them will not change appreciably for the duration of a data pulse. (In as much as individual conductors have an inductance, if only a small inductance is needed, the stray inductance of an array of conductors may be sufficient). Accordingly, if the totem pole cells are all in the same state, all high or all low, most of the time, and make only very brief pulses to the other state for data transfer, then the currents will remain generally unchanged.

[0037] When the totem pole cells are used in reverse just to control the voltage on V_{cc} , the various totem pole cells will change state so as to modulate the current into the storage capacitors on V_{cc} to maintain the charge nearly constant. Without regard to the ballast block, the duty cycle of individual totem pole inputs would be unimportant as long as the overall average duty cycle was correct to achieve the instantaneous current required. This is not so using ballast blocks, as the average voltage on the several current paths must be equal over a time that is short

compared to the time constant of the inductors. Thus the various totem pole cells should be operated with equal duty cycle, but the duty cycles are preferably multi-phased and interleaved to minimize the ripple voltage on V_{cc} .

[0038] When used to transfer data as well as to control the voltage V_{cc} , obviously the totem pole cells cannot all be in the same state all of the time and still output useful data. None the less, the currents will remain fairly well matched if the duty cycle of each line is the same over time. That is, if, over a time that is short compared to the time constant of the ballast block, the average voltage on each line is the same. If every line spends x percentage of the time low and $(100 - x)$ percentage of the time high, then the currents will remain matched. There are several ways that this can be accomplished.

[0039] One is if each data word is preceded (or followed) by its complement, having the same timing pulse width, then the word and its complement will average out.

[0040] Another scheme would be to encode the data using a modified "NRZ" data stream. NRZ data is usually used with serial data buses that are transformer coupled, to ensure equal polarity over time so that the transformer excitation

does not have a dc component and saturate the core. The same considerations can be applied to the ballast block 105, except that the objective is to achieve an equal average voltage, not zero average voltage.

[0041] In this invention, it is contemplated that a primary function of the totem pole cells is to control the voltage on Vcc, and that may be the only function in many applications. However, given the large number of input pins devoted to supplying power to a typical processor, there may be an advantage to using these input pins for the dual purpose of supplying power and transferring data. As an example, not a limitation, the power input pins could, using the teachings above, provide data and control to the power converter itself or to other slower peripherals within a computer. The higher voltage on such "data" would have a higher signal to noise ratio, and their lower impedance may be able to drive longer lines and reduce susceptibility to noise.

[0042] This is not to lose track of the advantage that may be enjoyed by injecting current into every data output line. Such a configuration may contribute significantly to the current into the Vcc of the processor, reducing the reliance on the primary power source. It also has the advantage of having

a higher voltage for the data signals, which would result in a higher signal to noise ratio. It may also provide a better drive to other circuits that may not operate at as low a voltage as the processor may.

[0043] Figure 12 shows a modified NRZ waveform. This waveform also presents a constant average voltage to the data lines, and the currents would remain matched through the ballast block. Figure 12a shows a "0", and figure 12b shows a "1". This modified NRZ waveform would yield an average 50% voltage and average 50% current, so the totem pole cells would have to control the current by varying their duty cycles (but with each having the same duty cycle) between data transfers.

[0044] If the data bus is a combined power control and data bus, then the NRZ waveform can be modified to have a duty cycle which can be varied over a wide range while still preserving the integrity of the data. Figures 13a and 13b show modified NRZ waveforms having a low duty cycle, and figures 14a and 14b show modified NRZ waveforms having a high duty cycle.

[0045] If a plurality of totem pole cells is used to control power, there will be a constant circulating current. This means that there is a constant loss in the die, either in the upper

MOSFETs as current is directed to V_{cc} or in the lower MOSFETs as the current is circulated back to the return. Whether the benefits of this method of power control are worth the penalty of this constant loss is a trade off of the system. It may very well be worthwhile, for the following reasons: First is the advantage of having constant currents in the power distribution bus. With constant currents, there is no di/dt in the power distribution bus, so the loop inductance is no longer a factor. The distribution bus can be quite long and the prime power source can be far away from the processor. With no di/dt , the near field is constant, so there is very little EMI generated by the distribution bus.

[0046] Second is the dynamic response. The currents into the storage can change from zero to full load and back just as quickly as the MOSFETs can be turned on and turned off, an extremely short time. This has a number of advantages. One is that the capacitors can be much smaller, as the current into the capacitor can be brought into compliance with the load current out of the capacitor very quickly.

[0047] A third advantage is that the voltage control can be very precise, with very small errors, even during large current

transients because errors can be detected and corrected very quickly. With more precise voltage control, the processor can be operated at a lower voltage for a significant savings in power. For example, if the MOSFETs have a 50 mv drop when on, the power loss will be about five percent of full load. If the processor can operate at 90 percent of the V_{cc} that it otherwise would, the power consumption is 81 percent, a net savings at full load of 14 percent. Unfortunately, the same five percent of full load power is present even at lower loads, but the break-even point would be about one third of full load and there would be a benefit at any greater load.

[0048] Fourth, with orders of magnitude less delay while current ramps up, then down, the total time that the current is above quiescent will be much shorter, for a significant savings in power as well as a having faster operational response to whatever data processing task required the increased current.

[0049] Fifth, with a plurality of constant current sources, the current can be switched to the part of the processor die where it is needed, reducing the consequences of unequal power load use on the die.

[0050] Sixth, the on resistance of MOSFETs that are part of the

processor die may be superior MOSFETs. In a MOSFET, the on resistance is a function of the number of cells and their voltage rating. The processor has a much better cell density than most power MOSFETs, and its voltage rating is particularly low. Accordingly, for a given area of silicon, the embedded MOSFETs are likely to have a much lower on-resistance than a comparable discrete or integrated circuit power MOSFET.

[0051] The plurality of constant currents need not be constant at all times. For some modes of operation, the currents can be reduced to a low value. If for some mode the currents can be reduced to 50 percent of the maximum, the power loss will be reduced by 75 percent. If the current can be reduced to 30 percent, for example, in a standby mode of operation, the power loss in the circulating currents will be reduced by 90 percent. Currents can be changed either by ramping down the constant current source or by switching currents in or out of the prime current source, the first method being simple but slow, and the second method being very fast, perhaps too fast. With reference to US Patent Application No to 10/708,523, "Buck (or Boost) Converter with Very Fast Output Current Transitions and Low Ripple Voltage" by using a dynamic induc-

tor output for the prime current source, the current can be transitioned very quickly while still having a defined di/dt determined by the circuit design. This patent application is owned by the same entity as the present invention, and it is incorporated herein by reference.

[0052] Further, if the processor has a known minimum current, one or more constant current sources totaling that minimum current can be connected directly to V_{cc} without switching. This directly connected minimum current may also be varied at the prime current source according to the mode of operation. Yet another choice is to have some portion of the current control on the processor die as embedded totem pole cells and have another portion located nearby (within and/or outside of the package). Current control of a portion of the current on the processor die would allow very fast voltage adjustment while most of the circulating current and its losses would be elsewhere.

[0053] There is a system trade off in the number of power input pins to use and how they are grouped. A reasonable number of constant current sources may be in the order of ten to twenty or somewhat more. On the other hand, a processor may have a large number of power input pins, and there is the possibility of terminating each with a totem

pole cell circuit for power control. The ten to twenty constant current sources can be divided using ballast blocks to get a multiple, say 80 to 150. For the pins that are wired through a ballast block from one constant current source, the average voltage on wires from the ballast block must be the same, which means that each pin must have the same duty cycle. This can be accomplished with a circulating algorithm. To envision how it works, consider a PWM pulse going into a shift register. All of the shift register outputs see the same pulse width (or duty cycle) but shifted in time. Using a shift register or a ring counter are possible control methods, but it is better to implement the function in logic so that the current can be adjusted more quickly if there is a sudden large change in demand. (In as much as a length of wire has inductance, it may be that the wire runs alone will have sufficient inductance. Such inductors are included when a ballast block is cited, even if there are no separate physical inductor components).

[0054] It is likely that no additional mask steps would be needed to implement totem pole cell power control. The power control inputs have the same MOSFET to V_{ss} and the same MOSFET to V_{cc} as ordinary data cells, except that the cur-

rent is higher. Just as power MOSFETs comprise a large number of cells, so would these MOSFETs, enough cells to handle the current with an acceptable voltage drop. The total number of cells would be a function of the total switched current, but how they are grouped would be a trade off of the specific design. One design might use ten inputs, and switch 10 amperes each, requiring a small number of MOSFETs, each comprising a very large number of cells. Another might use 80 inputs of 1.2 amperes each, requiring a larger number of much smaller MOSFETs. The total cells used would be about the same.

[0055] Figure 15 shows a switched current power converter 150 that is powering a processor 151. The processor 151 has a processor die 152, and it may be mounted on an interposer 153. The switched current power converter 150 may comprise a source of multiple constant currents 155 having a plurality of constant current outputs 163-163 and switched current switching means distributed across a voltage regulator module (VRM Pod) 154, the interposer 153, within the processor package 151 and on the processor die 152.

[0056] As an illustration, not a limitation, the source of multiple constant currents 155 may comprise a constant current

source 156 comprising an inductor 157, a switching means 156 and a catch diode means 159. One skilled in the art of power conversion will recognize the constant current source 156 as a usual buck converter, though it has no output capacitor and it is controlled by a controller (not shown) such that the current I in the inductor 157 is held approximately constant. The source of multiple constant currents 155 may further comprise a matrix transformer array 160 excited by the constant current I and switched by push-pull switching means 161 and 162 operated at "100 % duty cycle".

[0057] If the secondaries of the matrix transformer array 160 are full wave rectified, the result is an array of constant current sources 163-163 each of which has a constant current output of value I (assuming a one to one ratio in each section of the matrix transformer). The use of a buck converter driving a matrix transformer array is an example, not a limitation. Any plurality of constant current sources may be used with this invention with equivalent results. "Constant current", as used in the specification and the claims, is not an absolute; it allows some variation within the error and control band of the current regulating scheme used as well as different levels of "constant cur-

rent" for different modes of operation of the processor.

[0058] First, please note the two paralleled and unswitched constant current sources that are connected in a connection point 168 to Vcc. These are connected directly to the processor and the several storage capacitors 164, 165 and 166 in the VRM Pod 154, on the interposer 153 and within the processor package 153. This provides an unvarying minimum current to the processor and is a preferred configuration if the processor has a known minimum operating current regardless of its state or operating mode. This is often the case with processors, and by not switching this current, there are no switch losses associated with the minimum current.

[0059] The next two of the constant current sources 163-163 are connected to switching means 170 and 171 on the processor die 152. There may be any number of such switching means on the processor die, and the illustration of just two is to keep the drawing simple while still showing the principle of the invention. Their function and operation are as described in the discussion of figure 6 above.

[0060] The next two of the constant current sources 163-163 are connected to switching means 172 and 173 that are not on the processor die 152 but that are located within the

processor package 151. The next two of the constant current sources 163–163 are connected to switching means 174 and 175 that are located outside of the processor package 151 and on the interposer 153. Finally the last two of the constant current sources 163–163 are connected to two switching means 177 and 178 that are located in the VRM Pod 154. In a practical system for controlling the voltage V_{cc} for a processor 151 the location of the various switching means 172 through 178 would be a design trade off. In present processor power supplies, the voltage control is usually entirely within the VRM Pod 154 or its functional equivalent. However, one of the teachings of this invention is the use of switched current power control, and that the closer to the processor die the switching means are, the less will be the effects of rapid changes of processor load current (di/dt) due to the stray inductance of the power distribution.

[0061] Note that there is an additional "outside switching means" 176 to return on the interposer 153. In some circumstances, it may be undesirable to have a current circulating through the return of the processor continuously even though the current is not needed to sustain V_{cc} . In such circumstances, the outside switching means 176 may be

turned on and the switching means 170 may be turned off. There may be a plurality of outside switching means configured in this manner, and the illustration of only one is to keep the drawing simple. The outside switching means may comprise a totem pole driver as well, with a switching means to Vcc. However it is contemplated that the usual application for an outside switching means would be for modes of operation where the current from the attached constant current source will not be needed and the outside switching means is turned on to reduce circulating current and its attendant power loss within the processor (or other integrated circuit) With reference to US. Patent No. 6,121,761, "Fast Transition Power Supply", a switched-charge voltage step change capability can be added, as represented by a switched charge switching means 179 and a charge transfer capacitor 180 in figure 15. It is preferred to use external MOSFETs, as the peak current is quite large. Also, the capacitors are very likely external to the processor, and the charge should be switched to the capacitor directly and not put through the processor interconnections. (If the capacitors were entirely within the processor package, that would be a different situation. Then the charge switching should be done

there).

[0062] There are many tradeoffs in the method of control. The preferred method relies on the processor itself to control the switched-current totem pole cells, both in response to errors and in anticipation of load changes. To represent this, figure 15 shows a data bus 169 from the processor 152 to the various switches of the circuit as well as to the constant current buck regulator 156. While shown as a separate data buss for illustration and explanation, it could represent data transferred back on the constant current lines as described above.

[0063] To sense errors, a flash a/d converter may measure the voltage V_{cc} of the processor 152 and provide a digital input of an error voltage to the processor 151, and the processor may make compensating changes in the switching means to control the voltage V_{cc} . This could be accomplished with any of a number of processing algorithms, but a discussion of some considerations of this process follow.

[0064] If the total capacitance is known, a change in voltage from a previous measurement and a time between the successive measurements may be used to calculate a direct quantifier of the current error, so a correct number of

switches can be changed quickly to reverse a voltage drift. Further, in anticipation of an operation, and knowledge of the estimated current required for the operation, the processor can command a change of current at the start of the operation and again at the end. If the estimate is a little off, the error may be corrected by the error algorithm, but this will eliminate the chance of large current mismatches, improving voltage regulation and reducing the size of the capacitor needed.

[0065] To show the benefit of having the processor anticipate the load and command a current at the same instant, let's first review the capacitor needed without such control. Assume that the idle current is 30 amperes, and the full load current is 130 amperes. Further assume that the voltage is sampled at 50 ns intervals and that it takes another 50 ns propagation delay to effect a change in current. The current could have an error of 100 amperes for 100 ns before it is corrected. To maintain the voltage within 50 mv, the capacitor would have to be 200 uf. On the other hand, if the processor anticipated and programmed the current to within 10 amperes with precise timing, the maximum error would be 10 amperes for 100 ns, and a 20 uf capacitor would suffice. (This is not the only factor to be considered

in sizing the capacitor, so it is likely to be larger regardless).

[0066] There is an additional problem in the open loop scenario. At the end of the 100 ns, there is an error voltage of 50 mv. Bringing the converter current into conformance with the load current stops the error from getting worse, but to correct the error would require as much again current to be applied for the next 100 ns. Unless the current capacity of the converter is oversized, that much current may not be available, so it may be necessary just to keep the voltage error from getting worse, but not corrected. All is not lost, though, because at some time the current will change again and the error will be in the other direction, canceling the original error voltage. This suggests that the voltage should be controlled to be slightly high at low load, and slightly low for high load. However, this is a usual characteristic of the voltage and current relationship specified for processors.

[0067] With the processor anticipating and commanding current changes in real time, the errors will never get as big, resulting in much better voltage regulation.

[0068] If it is desired not to involve the processor in the voltage control, a power control co-processor can be used, either

embedded on the processor die or separate. A degree of communication for operational and mode changes is desirable none-the-less. Changes in the current source currents at the prime power source require a communication link to that circuit as well.